

Stelvio


Introduction

Stelvio is a program devoted to solving orthodox SPGs (Shortest Proofgames). To make things bit clearer, its easiest to consider an example. Let us look at a masterpiece from Finland.

Unto Heinonen

Probleemblad 1998

1. Prize



SPG in 23.0 moves

This can be checked by Stelvio, and in less than 1 minute (depending on hardware and other parameters). Stelvio comes to the conclusion, that there is exactly one way to reach this position in 23.0 moves:



The successful strategy and the unique solution can subsequently be displayed:



The name

Stelvio, with its 2757m of altitude above sea level, is one of the highest and one of the most beautiful mountain passes of the Alps. As a bike rider, I've been up there many times and have good memories of it. I needed a name, and I liked the way Stelvio sounds. This name breaks the

"tradition" of naming the program after a famous mathematician like Jacobi or Euclide. Let's say the idea of naming it Scholze was not very appealing...



There is another reason why Stelvio is a good name, as I found out by coincidence in fall of 2024: The anagram "Solve it"!

Compatibility

As Stelvio is written in Java, it can run on various platforms, among them Windows, Mac and Linux.

Requirements

Solving SPGs is often memory-intensive, especially for SPGs with a lot of free moves. For that reason, it can be very beneficial to run Stelvio with most of the RAM that you have on your machine. You can find the amount of RAM you have in the System properties, common values for a notebook would be 8g-32g. I'd say 4g is bare minimum (can Tetris run with less?).

Btw: Running Stelvio with way more RAM than is needed for a particular SPG will likely increase the solving time for this SPG, as there are costs involved accessing large amounts of memory on hardware level. You can witness this e.g. in the required ramp-up time: When given a lot of memory, Stelvio needs much longer to get going. This is due to the fact that a large cache needs to be initialized. But for SPGs that take a long time to solve, and added 20 seconds or so up front do not matter. When giving Stelvio almost all or all of your RAM, it is advised to close other applications so that the memory is in fact free to use.

Installation

As Stelvio is written in Java, there needs to be a corresponding Java Runtime Environment installed for it to run. You can get your OS-specific version from here: <https://www.oracle.com/ch-de/java/technologies/downloads/>. I have tested Stelvio with Java 11 and Java 17. E.g. for Windows, you can download Java 17 with this link: https://download.oracle.com/java/17/archive/jdk-17.0.12_windows-x64_bin.exe. You then need to run the downloaded file and install it to a directory of your choice, recommended is "c:\java\java17" (this path should not contain spaces). Now download stelvio<version>.zip and unzip it to a directory of your choice, e.g. "c:\spg\stelvio". The zip file contains the following:

- bin folder containing some *.jar files (the code)
- doc folder containing PDFs (documentation)
- stelvio<version>.bat (the file to run)
- stelvio<version>.sh (the file to run on Linux)
- stelvioUI.ini (the parameter file)
- problems.txt (example input)

Adjust stelvio<version>.bat

You should open stelvio<version>.bat in a text editor and adjust the maximum memory settings. Per default, it says "-Xmx8g" in the file, which means Stelvio is allowed to use 8g of RAM. If you have more memory available, then give most of it to Stelvio, e.g. "-Xmx16g" for 16g of RAM. Also make sure that the path that points to the Java Runtime Environment is correct. The default is "c:\java\java17\", but you need to adjust this in case you opted to install Java at some other location.

To test your installation, just double-click stelvio<version>.bat, which should run Stelvio with the provided problems.txt file.

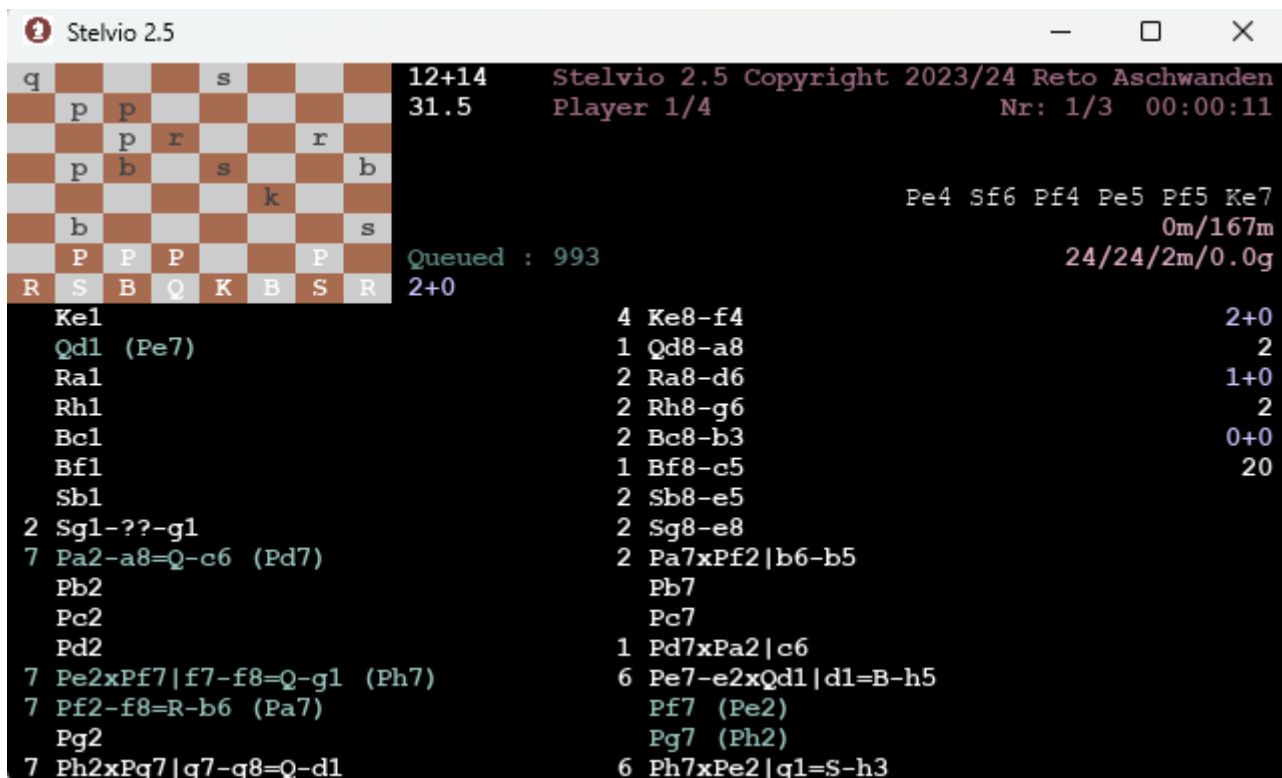
Troubleshooting

If Stelvio does not start when double-clicking on the stelvio<version>.bat file, then most likely the Java Runtime Environment cannot be found or its version is too old. Check if the path contained in stelvio<version>.bat pointing to java.exe is correct and adjust if needed.

UI

The simple terminal-based UI gives some impression where Stelvio stands in the solving process.

Here is what is displayed for a strategy player:



What it all means:

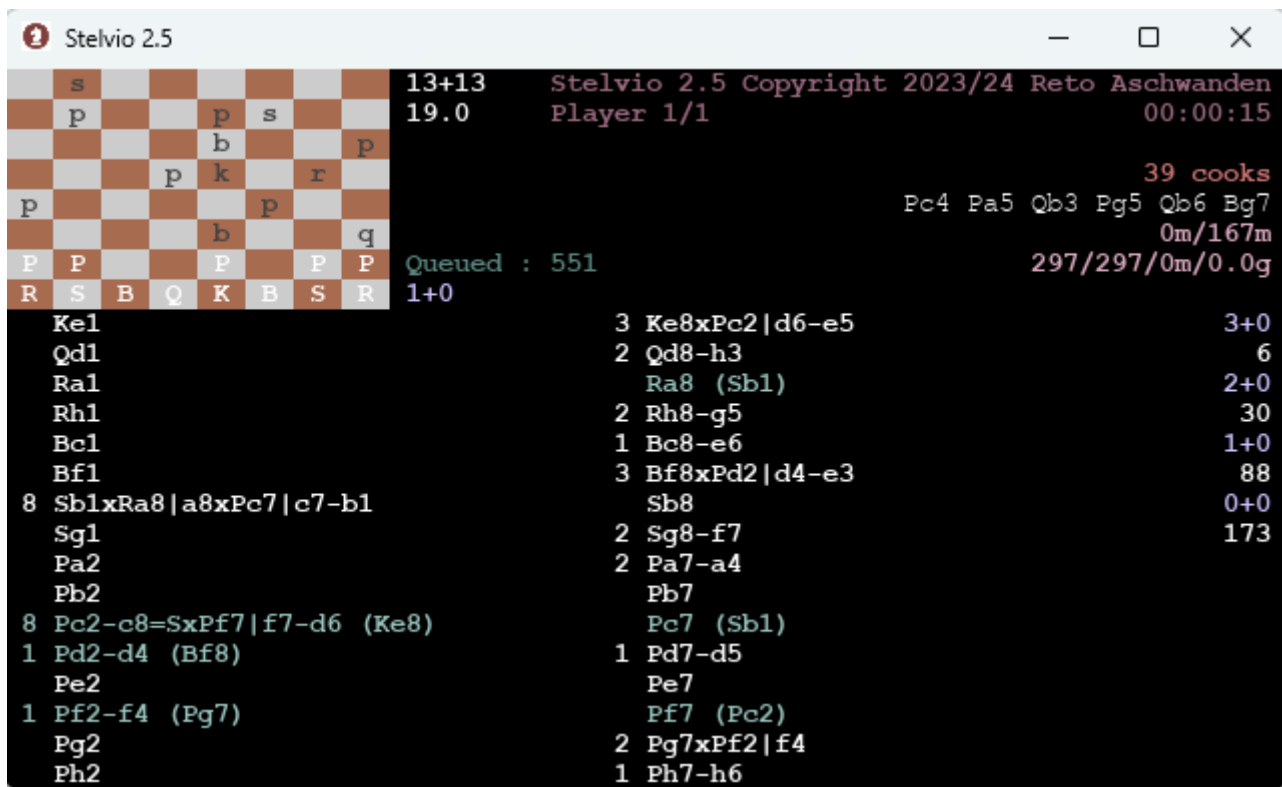
- Below the copyright notice:
 - Player 1/4: Which page is being displayed. We are currently looking at the page for strategy player 1 out of 4 such players. Pages can be switched using left-arrow/right-arrow in order to see pages for other strategy players/strategy seekers.
 - Nr: 1/3: The SPG being solved. In the example, there are 3 SPGs contained in the input file, and we are currently solving the first SPG.
- Next to the board:
 - Number of pieces and number of moves are obvious.
 - Queued: 993: Number of strategies that have been found that still need to be played. The numbers here are before any possible strategy splitting.
 - 2+0: Number of free moves for the current strategy by color, so 2 for white and 0 for black.
- On the right:
 - Solving time so far (hh:mm:ss).
 - Current move path (not to be taken too seriously, this can be stale/inconsistent data for technical reasons. If garbage is displayed, this has no effect on the solving itself, only the display is invalid).
 - Cache/strategy metrics:
 - 0m/167m: Currently, the playing cache is basically empty (0 million). Cache can hold 167 million positions (this is actually just an approximation, but good enough). This cache is bigger if you have more RAM available, which can be very useful depending on the SPG in question.
 - 24/24/2m/0.0g:

- 24: The number of strategies already taken off the strategy queue (before any strategy splitting).
 - 24: The strategy number being played (possibly after strategy splitting). In the current case, no strategy splitting was needed, and therefore the numbers coincide.
 - 2m: Number of moves played for the current strategy in millions.
 - 0.0g: Total number of moves played for all strategies so far, in billions. Currently, less than 100m have been played, so this displays 0.
- Below the board:
 - The current strategy, i.e. what each piece does, white on the left and black on the right, with associated move count per piece.
 - The column on the far right: Histogram of already played strategies (top 8 entries, as there is limited space). So in the example, 2 strategies with 2+0 free moves have already been played, next to strategies with 1+0 and 0+0 free moves.

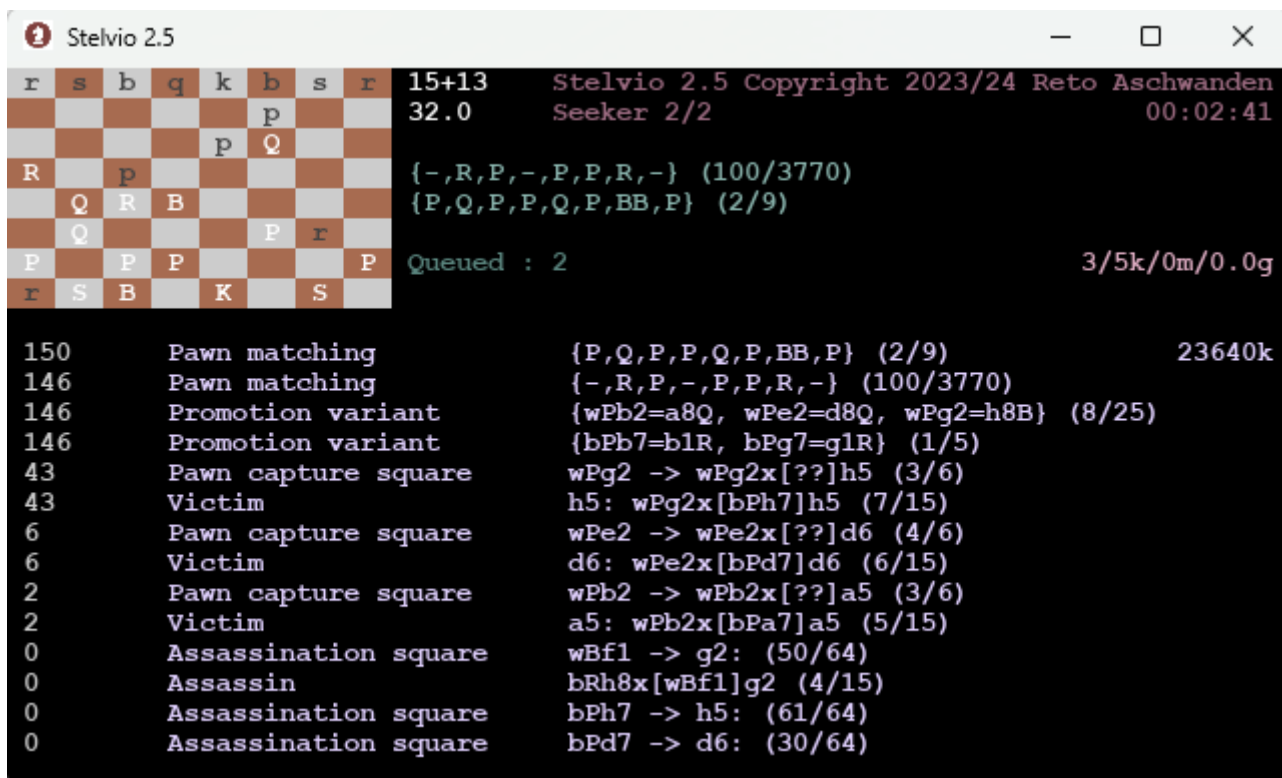
As already displayed above, when a solution is found, pressing up-arrow/down-arrow displays the solution and the successful strategy.

When there are cooks, Stelvio counts the number of strategies that contribute to cooks, in the below example 39 cook strategies were found so far. In case strategies which can be uniquely played are found, these are counted as valid solutions and displayed as well.

A brief note on what Stelvio counts as a cook/solution: In case a strategy can be played in multiple ways, then this counts as a cook strategy for Stelvio. This is the correct approach to take for 99.9% of the SPGs out there. In case a SPG requires an even number of half-moves, e.g. 22.0, then any solution/cook needs to have an even number of half-moves. Correspondingly for SPGs requiring an odd number of half-moves. This seems a good approach to take, since who is at play is part of the position to reach. So in case in the example one could reach the diagram position in 21.5 moves, then this does not count for Stelvio. Any shorter move-path with same parity of half-moves does count though, so in the example, Stelvio would count solutions/cooks in {21.0, 20.0, 19.0, ... 1.0} moves.



Above we have seen the UI for a strategy player, let's now look at what is displayed for a strategy seeker:



- Next to the board:
 - The green lines: How the pawns are assigned to partitions currently in the strategy seeking process. Partitions are the types of pieces on the board that are mutually exclusive, we have partition values (K, Q, R, WB, BB, S, P). WB stands for white square bishop and BB for black square bishop. Top line is black assignment {-,R,P,-,P,P,R,-}, lower line is white assignment {P,Q,P,P,Q,P,BB,P}. This means for white, that a/c/d/f/h pawns remain pawns, b/e-pawns are

queens in the diagram position and g-pawn is a black square bishop in the diagram. For black, c/e/f pawns remain pawns, the b/g-pawns are rooks visible in the diagram and a/d/h pawns are captured. At the outset of strategy seeking, all possible assignments of pawns to partitions are calculated, in the example we have 9 for white and 3770 for black. During strategy seeking, these assignments are all gone through in a big loop, and we are currently at assignment 2 for white and 100 for black. This information gives you a hint how far strategy seeking and therefore solving already is.

- Queued : 2. The number of strategies in the queue that are waiting to be played. This queue is the link between all the strategy seekers and all the strategy players.
- On the right:
 - Solving time so far (hh:mm:ss).
 - Strategy metrics:
 - 3/5k/0m/0.0g:
 - 3: The number of strategies that were retained after strategy analysis and which therefore need to be played.
 - 5k: The number of strategies initially found (before strategy analysis), in thousands.
 - 0m: Since the strategy seeker is not playing, this number remains 0.
 - 0.0g: Since the strategy seeker is not playing, this number remains 0.

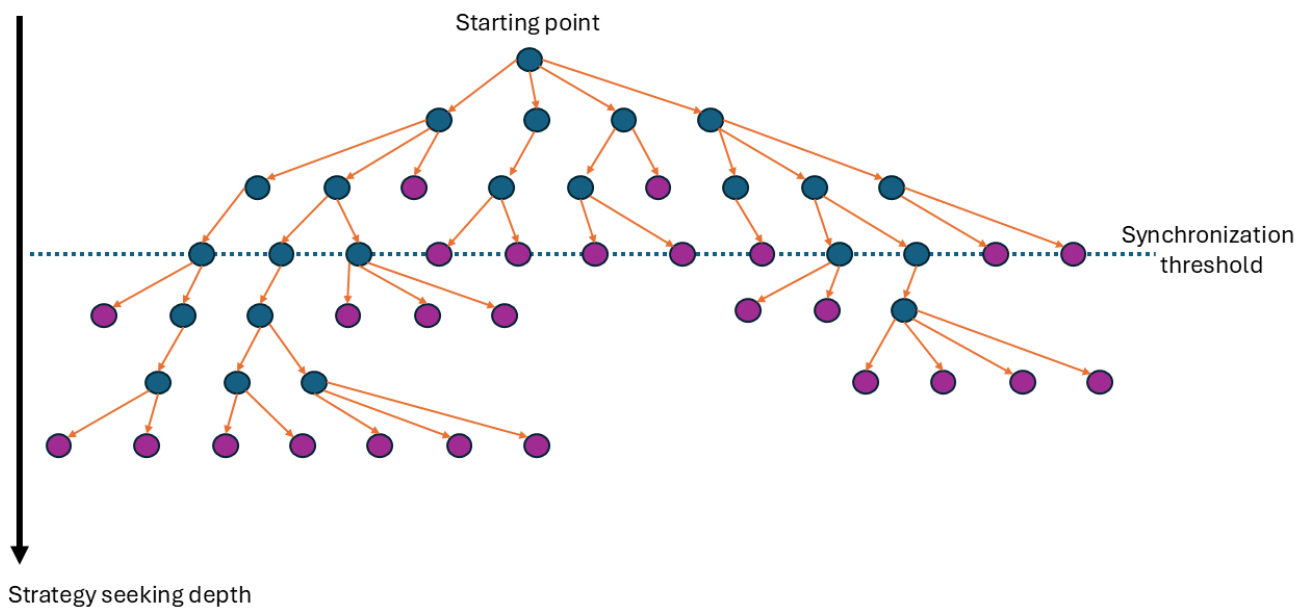
From version 1.6 onwards, the current strategy seeking path can be displayed. This is helpful in order to get an estimate of how long the seeking process will take. The first 14 levels of seeking are shown, together with how long the corresponding seek node has been active (in seconds, the number on the left). This information is somewhat analogous to the first 6 moves displayed when playing, only here, we have the first 14 "seeking moves". The information should be read top-down, so the two pawn matching choices have been active for 150 seconds respective 146 seconds. The next thing the seeking algorithm looks for are promotion variants. After that, a pawn capture square is chosen for the Pg2, currently we chose h5. Btw, the amount of possible options displayed to the right is sometimes only an estimate, but better an estimate than nothing.

If you are in histogram mode, then up-arrow/down-arrow switches between histogram and this new seeking information. Gathering this information takes a bit of time and in case you are not interested, it can be fully turned off. See the StelvioParameters file for details.

Strategy seeking parallelization algorithm

With version 3.0, Stelvio uses an embarrassingly simple but much improved approach to parallelize the strategy seeking workload. Once I saw this approach, I could not figure out why I had not seen it earlier.

On an abstract level, strategy seeking traverses a tree of search nodes, usually billions of nodes. At the top of the tree (i.e. the starting point), the strategy is basically empty, containing only the things that are visible from the diagram position. While moving down the search tree, the strategy becomes more and more defined, until the strategy is either fully defined or we have reached a dead end. All found strategies are leaves in this tree. The converse is not true: Not all leaves are strategies, as most search paths are simply dead ends that do not lead to a fully defined strategy.



For the new approach, there is a synchronization threshold defined by the parameter `strategySeekingSyncDepth` (dashed line in the image). This is the tree depth at which all the strategy seekers compete for tree nodes. All tree nodes above this level are computed by all strategy seekers. As the tree grows exponentially, the top few levels pale in size compared with the lower levels, so computing the top levels multiple times is usually negligible. When a seeker reaches the synchronization threshold level, it tries to acquire the corresponding node n . Two things can happen:

- Acquiring the node n is successful (for any node on the synchronization level, this is true for exactly one seeker). This seeker is then solely responsible to compute the whole subtree of n . No parallelization takes place for this subtree.
- Acquiring the node is not successful. The seeker interprets this path as a dead end and goes on searching (moving backwards first)

All the strategy seekers agree on the order of all the nodes that are found up to the synchronization level (i.e. the subtree defined by the synchronization level and all nodes above it is traversed in deterministic fashion by all seekers). So all these nodes can be viewed as numbered, and all seekers agree on these node-numbers. The synchronization mechanism simply keeps track of the maximal node-number max_n on the synchronization level that has been acquired. An acquisition attempt for node-number n is successful if and only if $n > \text{max_n}$. Such a successful attempt then increases max_n to n . Therefore, any later seeker trying to acquire this same node n fails, because $n \leq \text{max_n}$.

On a side note: The invariant that all seekers must find the same amount of nodes above the synchronization level is checked at the end of strategy seeking. If the numbers do not coincide, then an internal error occurs.

Effects of `strategySeekingSyncDepth` (i.e. synchronization level)

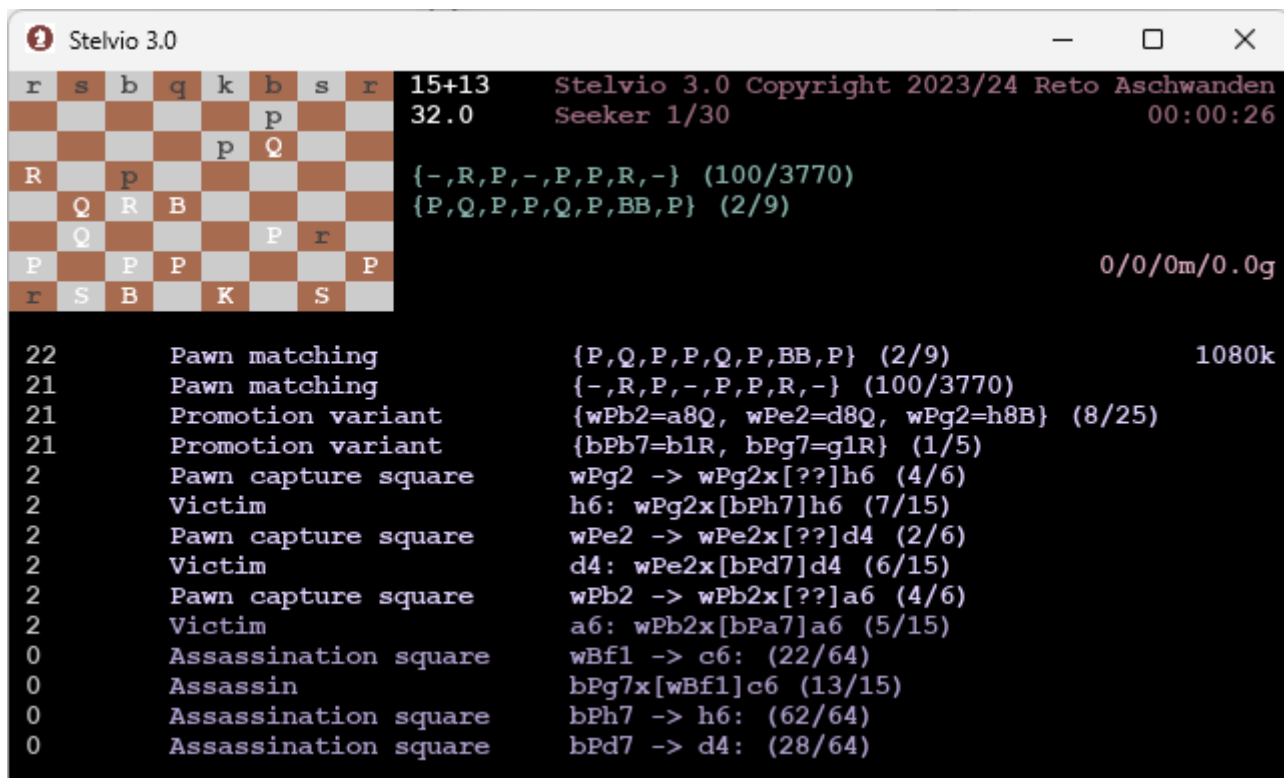
The default level is 12 and should be ok for most SPGs. I assume that levels between 8 and 18 could

be useful.

- Level too low: This can lead to strategy seeker starvation, as the work items get too big in a skewed search tree. So you might have 10 seekers, but after a short time, only 1-2 are still active. That is a sign that you should increase `strategySeekingSyncDepth`.
- Level too high: If the level is set too high, then the part of the tree that all seekers calculate becomes large (if you set the level to say 100 (i.e. a level which is greater than the tree depth), then all seekers calculate the whole tree and nothing is parallelized). So arbitrary high values are not useful either.

So there is a tradeoff to be made here and the optimal value depends on the SPG at hand.

The synchronization level is actually visible in the UI:



The current strategy seeker path can be seen here (top-down). So first, the pawn partition matchings are chosen, then promotion variants etc. The lower 4 nodes are displayed in slightly darker violet. These darker nodes depict nodes below the synchronization level, and they are therefore not parallelized. Btw: Some nodes do not have a useful UI representation (sometimes). So even if you have `strategySeekingSyncDepth=12`, you might not see all 12 nodes in the UI (like in the example: there are only 9 lighter nodes displayed).

Histogram mode


In order to get a first impression if an SPG can be solved in a reasonable amount of time, it can be useful to simply search for all the strategies without actually playing them. This is what the histogram mode is for. In this mode, strategies are added up and grouped by white/black free moves.



So in the example, there are 72 strategies currently found by this strategy seeker (numbers are per seeker) with 2+0 free moves, 138 with 1+0 and 403 with 0+0 free moves. When all seekers are finished, the strategy numbers are added up and written into the output file. It is also possible to attain a partial histogram: You can tell Stelvio to only start adding up strategies after strategy number X. Beware though that strategy numbers are only deterministic in case only 1 strategy seeker is used. This can be useful if you want to know what is left in terms of strategies, if Stelvio already solved up to strategy X.

Last move retraction

In case a king is in check in the diagram, it can be very helpful to retract the last half move and solve all implied shorter SPGs and combine their results thereafter. By default, this is what Stelvio now does in such a case. While solving, the retracted move is displayed next to the diagram, as well as how many such moves there are. In the example, there are four possible last moves, and we have currently retracted the second. The problem displayed is the one with the move retracted. In the special e.p. retracted case, 2 half moves are retracted, as these 2 half moves cannot be split up. In the special castling retracted case, Stelvio makes sure that the king and involved rook are not moving beforehand. In histogram mode, the histogram of all sub-SPGs are written to the output file, as well as an overall histogram. Beware that this overall histogram will not coincide with the histogram you get when solving the SPG without retraction. This retraction functionality can be switched off by a new parameter named `retractionMode`. Limitation: Last move retraction and save/read strategies from/to file are not currently supported in combination. Btw: The `retractionMode=always` (i.e. even if king is not in check) will probably be added in a future release.



Some thoughts on parallel solving

Both histogram and normal solving are executed in parallel. Modern computers come with several CPUs, so calculating in parallel can be beneficial for performance.

The basic architecture is quite simple: There are x strategy seekers and y strategy players. The numbers x and y can be parametrized to suit the SPG at hand and your hardware. When running in parallel, the strategy seekers search for strategies and put the found strategies into a queue. The strategy players on the other hand consume strategies from this queue and try to play them. The granularity for the strategy seekers is a single pawn partition matching combination plus some further few seeking choices. This can be a bit coarse, and I might adjust this in the future. In the example below, there are $9 * 3770$ such combinations, which the seekers consume one by one until none are left.

```

Stelvio 2.5
15+13  Stelvio 2.5 Copyright 2023/24 Reto Aschwenden
32.0   Seeker 2/2 00:02:41

  r s b q k b s r 15+13  Stelvio 2.5 Copyright 2023/24 Reto Aschwenden
  p 32.0   Seeker 2/2 00:02:41
  R  { -,R,P,-,P,P,R,-} (100/3770)
  Q  {P,Q,P,P,Q,P,BB,P} (2/9)
  Q
  P P P P P
  r S B K S
Queued : 2 3/5k/0m/0.0g

150 Pawn matching {P,Q,P,P,Q,P,BB,P} (2/9) 23640k
146 Pawn matching {-,R,P,-,P,P,R,-} (100/3770)
146 Promotion variant {wPb2=a8Q, wPe2=d8Q, wPg2=h8B} (8/25)
146 Promotion variant {bPb7=b1R, bPg7=g1R} (1/5)
43 Pawn capture square wPg2 -> wPg2x[??]h5 (3/6)
43 Victim h5: wPg2x[bPh7]h5 (7/15)
6 Pawn capture square wPe2 -> wPe2x[??]d6 (4/6)
6 Victim d6: wPe2x[bPd7]d6 (6/15)
2 Pawn capture square wPb2 -> wPb2x[??]a5 (3/6)
2 Victim a5: wPb2x[bPa7]a5 (5/15)
0 Assassination square wBf1 -> g2: (50/64)
0 Assassin bRh8x[wBf1]g2 (4/15)
0 Assassination square bPh7 -> h5: (61/64)
0 Assassination square bPd7 -> d6: (30/64)

```

Parallel solving examples

The screenshots below are from version 1.27, but the insights gained are still valid in the current version.

Same SPG but different concurrency metrics

I give a few examples of the effect of parallelizing, when it is beneficial and when it is not. Let's take the 2023 Andernach TT winner by Michel. In non-parallel mode (the only mode available in Stelvio 1.2), this SPG is solved on my notebook in some 35 minutes:

```

Stelvio 1.27
12+13  Stelvio 1.27 Copyright 2023 Reto Aschwenden
16.5   00:35:39

  r s  q  s
  p  p p p R
  k  p
  b r
  P P P P P P P
  S B Q K R
2+0
2 Ke1-??-e1 4 Ke8-f5 6+0
  Qd1 (Rh8) Qd8 4
4 Ra1-g7 Ra8 5+1
  Rh1 7 Rh8xPf2|f2xBf1|f1xQd1|d1xSg1|f3 1
  Bc1 Bc8 (Pa2) 5+0
  Bf1 (Rh8) 2 Bf8-e3 53
  Sb1 Sb8 4+1
1 Sg1-f3 (Rh8) Sg8 28
8 Pa2xPb7|b7xBc8|c8=QxPc7|c5-d1 Pa7 3+2
  Pb2 Pb7 (Pa2) 40
  Pc2 1 Pc7-c5 (Pa2) 2+3
  Pd2 Pd7 21
  Pe2 Pe7 1+4
  Pf2 (Rh8) Pf7 5
  Pg2 1 Pg7-g6 4+0
  Ph2 1 Ph7-h5 730

```

Both strategy seeking and strategy playing take significant time for this SPG, so splitting up is beneficial. With 1 seeker and 1 player, the solving time drops to 27 minutes in Stelvio 1.3:

Stelvio 1.27

12+13 Stelvio 1.27 Copyright 2023 Reto Aschwanden
16.5 Player 1/1 00:27:18

{S,P,P,P,P,Q,P,P} (14/14)
{P,R,-,P,P,P,P,P} (3/3)
{Queued: 0}

Done. Found 1 solution
478184/478k/0m/4.4g

2+0

2 Ke1-??-e1	4 Ke8-f5	6+0
Qd1 (Rh8)	Qd8	4
4 Ra1-g7	Ra8	5+1
Rh1	7 Rh8xPf2 f2xBf1 f1xQd1 d1xSg1 f3	1
Bc1	Bc8 (Pa2)	5+0
Bf1 (Rh8)	2 Bf8-e3	53
Sb1	Sb8	4+1
1 Sg1-f3 (Rh8)	Sg8	28
8 Pa2xPb7 b7xBc8 c8=QxPc7 c5-d1	Pa7	3+2
Pb2	Pb7 (Pa2)	40
Pc2	1 Pc7-c5 (Pa2)	2+3
Pd2	Pd7	21
Pe2	Pe7	1+4
Pf2 (Rh8)	Pf7	5
Pg2	1 Pg7-g6	4+0
Ph2	1 Ph7-h5	730

As no strategy requires a lot of memory, we can also split up strategy playing without harm. With 1 seeker and 2 players, the solving time drops to 20 minutes:

Stelvio 1.27

12+13 Stelvio 1.27 Copyright 2023 Reto Aschwanden
16.5 Player 2/2 00:20:01

{S,P,P,P,P,Q,P,P} (14/14)
{P,R,-,P,P,P,P,P} (3/3)
{Queued: 0}

Done. Found 1 solution
232959/232k/0m/2.2g

0+0

2 Ke1-??-e1	4 Ke8-f5	6+0
Qd1	Qd8	3
2 Ra1-f4 (Rh8)	Ra8	5+1
Rh1	5 Rh8xBf1 f1xRa1 f4xSg1 f3	1
Bc1	Bc8 (Pa2)	5+0
Bf1 (Rh8)	2 Bf8-e3	25
Sb1	Sb8	4+1
1 Sg1-f3 (Rh8)	2 Sg8xPf2 f6-g8	13
9 Pa2xPb7 b7xBc8 c8=RxPc7 c6-g7	Pa7	3+2
Pb2	Pb7 (Pa2)	21
Pc2	1 Pc7-c6 (Pa2)	2+3
Pd2	Pd7	12
Pe2	Pe7	1+4
3 Pf2-f6 (Sg8)	Pf7	3
Pg2	1 Pg7-g6	4+0
Ph2	1 Ph7-h5	362

In the above configuration, the queue is sometimes empty. That means that the players have nothing to do at times, so an additional seeker will likely reduce the solving time. With 2 seekers and 2 players, the solving time drops to 18 minutes:


```

Stelvio 1.27
12+13 Stelvio 1.27 Copyright 2023 Reto Aschwenden
16.5 Player 2/2 00:18:19

{S,P,P,P,P,Q,P,P} (14/14)
{P,R,-,P,P,P,P,P} (3/3)
{Queued: 0}
Done. Found 1 solution
240371/240k/0m/2.2g

0+1

2 Ke1-??-e1 4 Ke8-f5 6+0
Qd1 Qd8 4
4 Ra1-g7 Ra8 5+1
Rh1 5 Rh8xBf1|f1xPf2|f6xSg1|f3 1
Bc1 Bc8 (Pa2) 5+0
Bf1 (Rh8) 2 Bf8-e3 25
Sb1 2 Sb8xPa2|b8 4+1
1 Sg1-f3 (Rh8) Sg8 16
7 Pa2xPb7|b7xBc8|c8=QxPc7|c7-b8 (Sb8) Pa7 3+2
Pb2 Pb7 (Pa2) 18
Pc2 Pc7 (Pa2) 2+3
Pd2 Pd7 10
Pe2 Pe7 1+4
3 Pf2-f6 (Rh8) Pf7 3
Pg2 1 Pg7-g6 4+0
Ph2 1 Ph7-h5 367

```

Just for fun, I tried with 4 seekers and 4 players, which still shaves a minute of solving time:

```

Stelvio 1.27
12+13 Stelvio 1.27 Copyright 2023 Reto Aschwenden
16.5 Player 1/4 00:16:56

{S,P,P,P,P,Q,P,P} (14/14)
{P,R,-,P,P,P,P,P} (3/3)
{Queued: 0}
Done. Found 1 solution
118282/118k/0m/1.1g

1+0

2 Ke1-??-e1 4 Ke8-f5 6+0
Qd1 2 Qd8xPa2|c7-d8 1
4 Ra1-g7 Ra8 5+1
Rh1 6 Rh8xSg1|f3xBf1|f1xPf2|f6-f3 1
Bc1 Bc8 (Pa2) 5+0
Bf1 (Rh8) 2 Bf8-e3 15
Sb1 Sb8 4+1
1 Sg1-f3 (Rh8) Sg8 9
6 Pa2xPb7|b7xBc8|c8=RxPc7|c7 (Qd8) Pa7 3+2
Pb2 Pb7 (Pa2) 8
Pc2 Pc7 (Pa2) 2+3
Pd2 Pd7 7
Pe2 Pe7 1+4
3 Pf2-f6 (Rh8) Pf7 1
Pg2 1 Pg7-g6 4+0
Ph2 1 Ph7-h5 174

```

As my notebook now runs out of hardware, increasing the numbers further will have a negative effect, since in the end, all seekers/players compete for the same underlying hardware.

Different SPGs and the usefulness of parallelism for them

Histogram mode in parallel (strategy seeking in parallel)

With the new improved method of parallelizing introduced with version 3.0, strategy seeking

parallelization should be useful in all the cases that require substantial strategy seeking time. As a rule of thumb, the number of captures contributes heavily to strategy seeking time, but also the number of possible promotions and of course SPG length.

Solving multiple strategies in parallel

- Parallel solving of n strategies at a time can be detrimental under some circumstances. The reason is that each played strategy needs its own position cache, and this cache requires a lot of memory. So in case you have 2 players instead of 1, then each position cache is only half the size. Now in case while playing a strategy, the position cache is close to full, then solving performance for this strategy decreases dramatically. Therefore, it can be slower to play strategies in parallel than to play them serially. A case where this happens is the length record problem (P1407171), which by the way still cannot be solved.
- In the Andernach TT winner above, parallel solving is pretty useful, cutting the solving time in half.
- Parallel solving is the most useful in case strategy seeking takes up most of the time, and you have a lot of hardware (CPUs) that can be used for parallelizing. As there is almost no limit for number of seekers (given enough available CPUs), you can go all-in in that respect. An order of magnitude can be shaved of the solving time in that way for P1386153.

Solving a single strategy in parallel

- With v4.0, Stelvio is able to play a single strategy in parallel. This approach does not suffer from the downside (cache size decrease) mentioned above, so this approach is often preferable. The parallelization technique used is on an abstract level exactly the same as the one used for strategy seeking, see the chapter "Strategy seeking parallelization algorithm". At a certain half-move depth (given by the parameter "slavePlayerHalfMoveSyncDepth"), all slave players compete for the subtree beneath the current move path. Say the synchronization depth is 7, and a slave player played 1.Sc3 a5 2.Se4 a4 3.Sg3 a3 4.e4. At this point, this slave player tries to acquire the rights to play the subtree beneath 4.e4. In case no one has yet acquired this subtree, then this slave player's request succeeds, and he goes on playing this subtree. Otherwise 4.e4 is a dead end for this slave player (the subtree is skipped).
- Performance expectations: There is a synchronisation penalty to be paid when playing a single strategy in parallel. The shared position cache is highly frequented by all involved slave players and this is not for free. So performance improvements are not proportional to the number of slave players. On my machine, I found about 5fold performance increase with 10 slave players and about 10fold increase with 30 slave players. Your mileage may vary. When parallelizing this way, the cache should be split up into a decent number of cache splits, since these cache splits are the items that are synchronized on (cache splitting). So "positionCacheSplitExponent" should be set somewhere in the 8-15 range. Otherwise, there is too much contention and the cache becomes a bottleneck.
- This new mode of parallelizing is useful most of the time. I would not use it in the case strategies are numerous and are played very quickly anyway (the usual case for massacre style SPGs). There are slight setup-costs involved for playing in parallel, so in such a case, playing single-threaded might be faster (or playing multiple strategies in parallel).
- These 2 parallelization techniques can be combined. So you might play 2 strategies in parallel with 4 slave players each.

- UI: Each slave player gets its own page. E.g. in case 2 strategies are played in parallel with 4 slave players each, then we have players 1.1/1.2/1.3/1.4 and 2.1/2.2/2.3/2.4 respectively.

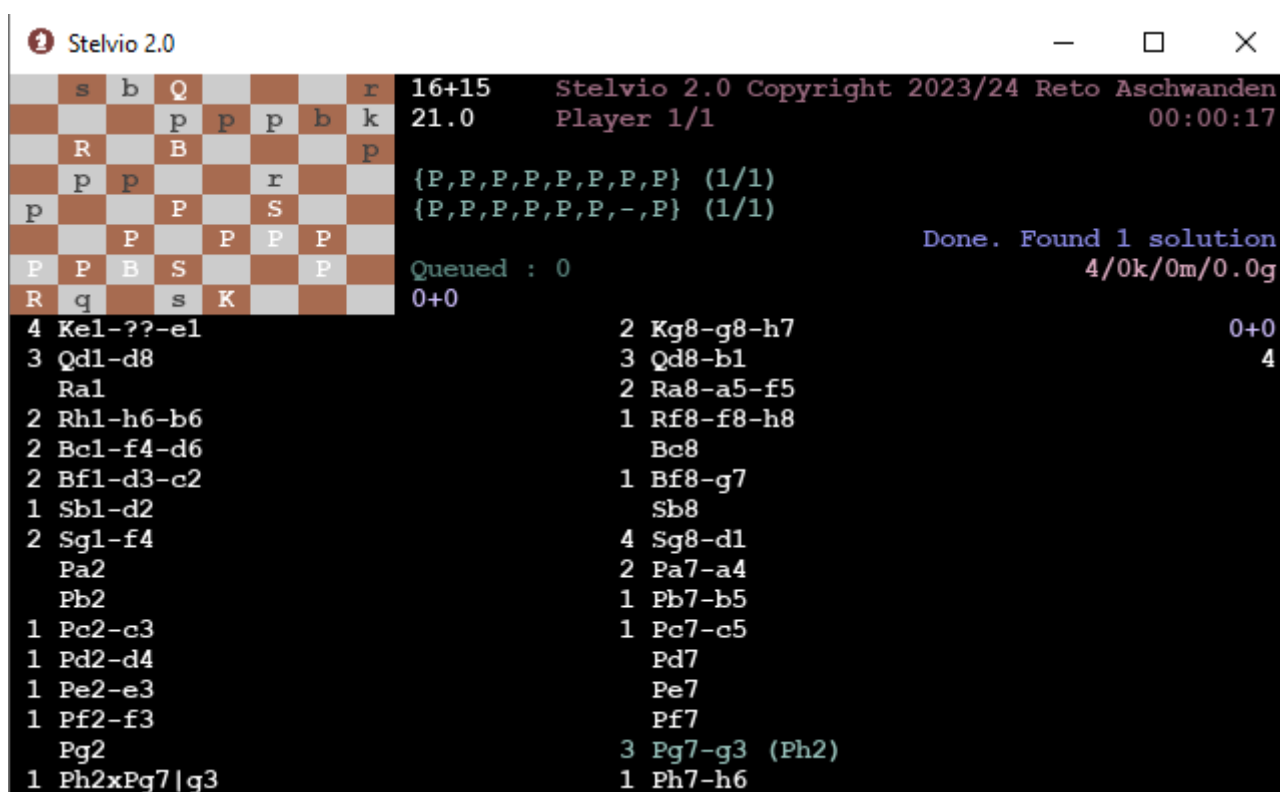
Current limitations, possible future developments

One big possible improvement comes to mind: Parallelize not only on one machine, but across a cluster of machines.

Collision detection modes

With v2.0, Stelvio has two modes of collision detection: A more involved mode (`expensiveCollisionDetectionMode = on`) and a less involved mode (`expensiveCollisionDetectionMode = off`). Both modes have their merit, as both can be orders of magnitude faster than the other, depending on the problem. In case `expensiveCollisionDetectionMode = default`, then the expensive mode is chosen for an SPG if and only if it has at least 28 pieces, which seems like a reasonable heuristic.

Collision detection is done for every found strategy. In case there are only few strategies to analyze, then the additional costs of more involved collision analysis will be negligible. One example where the involved analysis is very beneficial is the 4th Prize in the Champagne Tourney 2023 by Peter van den Heuvel.



The screenshot shows the Stelvio 2.0 interface. On the left is a chessboard with pieces: White (P, P, P, P, P, P, P, P) and Black (R, B, S, P, P, P, P, P). The board is labeled with files a-h and ranks 1-8. Below the board is a list of moves for both sides. The right side of the interface shows the game status: 16+15, 21.0, Stelvio 2.0 Copyright 2023/24 Reto Aschwandten, Player 1/1, 00:00:17. It also shows the search progress: {P,P,P,P,P,P,P,P} (1/1), {P,P,P,P,P,P,-,P} (1/1), Done. Found 1 solution, 4/0k/0m/0.0g, Queued : 0, 0+0.

16+15 21.0 Stelvio 2.0 Copyright 2023/24 Reto Aschwandten
Player 1/1 00:00:17

{P,P,P,P,P,P,P,P} (1/1)
{P,P,P,P,P,P,-,P} (1/1)

Done. Found 1 solution
4/0k/0m/0.0g

Queued : 0
0+0

4 Ke1-??-e1 2 Kg8-g8-h7 0+0
3 Qd1-d8 3 Qd8-b1 4
Ra1 2 Ra8-a5-f5
2 Rh1-h6-b6 1 Rf8-f8-h8
2 Bc1-f4-d6 Bc8
2 Bf1-d3-c2 1 Bf8-g7
1 Sb1-d2 Sb8
2 Sg1-f4 4 Sg8-d1
Pa2 2 Pa7-a4
Pb2 1 Pb7-b5
1 Pc2-c3 1 Pc7-c5
1 Pd2-d4 Pd7
1 Pe2-e3 Pe7
1 Pf2-f3 Pf7
Pg2 3 Pg7-g3 (Ph2)
1 Ph2xPg7|g3 1 Ph7-h6

The involved analysis detects that the black queen must pass through e1, eliminating any free moves that white would otherwise have. In case this is solved with `expensiveCollisionDetectionMode = off`, then solving takes hours if not days.

Any massacre style SPG will be solved quicker with `expensiveCollisionDetectionMode = off`, as there are billions of strategies to analyze and involved analysis is not needed. But not only massacres can benefit from less collision detection, the following SPG (P1013138) by me, Michel and Gerd is also

solved quicker with less analysis (by a factor 4).

Stelvio 2.0								Stelvio 2.0 Copyright 2023/24 Reto Aschwanden							
13+11								Player 1/1							
18.5								00:09:21							
{P,P,P,P,P,-,-,P} (10/10)								Done. Found 1 solution							
{S,P,P,S,P,Q,P,-}								7152/77k/0m/0.0g							
Queued : 0								0+1							
2 Kg1-g1-h1								Ke8							
2 Qd1-h3								Qd8							
1 Ra1-g1								Ra8							
4 Rf1-f1xPd7 c1xPh7 h5-h7								Rh8							
1 Bc1-e3								2 Bc8-b5 (Pc2)							
Bf1 (Pf7)								Bf8							
2 Sb1xPf7 f1								Sb8							
1 Sg1-f3								Sg8							
Pa2								1 Pa7-a6 (Pc2)							
Pb2								Pb7							
4 Pc2xBc8 b5xPa7 a6-a7								Pc7							
1 Pd2-d4								7 Pd7xPe2 e2-e1=S-c1 (Rh1)							
Pe2 (Pd7)								Pe7							
1 Pf2-f4								6 Pf7xPg2 g2-g1=RxBf1 f1 (Sb1)							
Pg2 (Pf7)								Pg7							
Ph2								1 Ph7-h5 (Rh1)							

Check protection

With v2.0, Stelvio can sometimes determine that check protection is needed. This can greatly reduce solving times. An example is the classic 1999 Phénix first prize by Michel (P1000010), which can now be solved in reasonable time on an ordinary machine.

Stelvio 2.5								Stelvio 2.5 Copyright 2023/24 Reto Aschwanden							
15+13								Player 1/1							
25.5								00:00:54							
Pa4 Sc6 Ra3 Se5 Rc3 Pb5								54m/167m							
Queued : 27								3/7/243m/0.2g							
7+0															
Ke1								3 Kc8-c8-a8							
2 Qd1-f6 (Pg7)								2 Qd8-e3 (Pd2)							
6 Ra1-b7-a1								3 Rd8-d8-f4 (Pg2)							
Rh1								2 Rh8-g8-g2							
Bc1								3 Bc8-h5							
3 Bf1-a6								3 Bf8-h2							
Sb1								3 Sb8-g4							
Sg1								2 Sg8-??-g8							
2 Pa2xPb7 b5								1 Pa7-a5							
Pb2								1 Pb7-b5 (Pa2)							
Pc2								1 Pc7-c6							
2 Pd2xQd8 e3-e4								Pd7							
Pe2								Pe7							
Pf2								Pf7							
3 Pg2xRa8 f4-f5								1 Pg7xQd1 f6							
1 Ph2-h4								Ph7							

By looking at the displayed strategy, it becomes apparent that check protection necessity for the

black king was detected, as a white rook is planned to make a detour (Ra1-b7-a1). Check protection gives rise to strategy splitting: Once check protection necessity is detected, all the possible ways to provide it are calculated. The original strategy is subsequently split into several sub-strategies. This splitting is visible in the UI in the metrics on the right hand side: In the example, we already took 3 strategies from the queue, but we are currently playing the 7th (sub-)strategy.

Btw: Strategy splitting can also arise from baseline collisions when it is unclear who makes a detour, e.g. Ra1-g1, Rh1-b1 is split into 2 sub-strategies.

Distribute free moves to sub-strategies

As of v4.3, it is possible to split up strategies further into many sub-strategies, which can be very useful in certain scenarios. This option is color specific and can be done for both colors separately, let us consider an example for white. Take a strategy which has white non-trivial paths Sb1x[Pb7]b5-b6 (5), Qd1-e2 (1), Pe2-e4 (1) and a total of 9 moves, so 2 free moves.

Below are all the ways the 2 extra moves could be consumed by white:

- Switchback by Sg1
- Switchback by Ke1
- Switchback by Bf1
- Qd1-e2 (3)
- Qd1-e2 (2), Pe2-e3-e4
- Sb1x[bPb7]b5-b6 (7) with 4 moves for first segment
- Sb1x[bPb7]b5-b6 (7) with 5 moves for second segment

You get the idea. So instead of playing the original strategy with 2 free moves, we could play these 7 sub-strategies as defined above. In many cases, playing all sub-strategies is slower, especially in case the original strategy does not take all that long to play anyway (there are fixed costs to be paid per (sub-)strategy). But the more free moves there are and the longer the strategy takes to play, the more effective this splitting option becomes. This is especially true in case playing the original strategy saturates the position cache, at which point solving performance falls off a cliff. Even in case all sub-strategy move trees combined have more nodes than the move tree of the original strategy, it might still be the case that the sub-strategies can be played faster, due to this cache-saturation effect.

As is obvious, the number of possible sub-strategies grows exponentially with the number of free moves, so splitting 10 free moves into sub-strategies might be a bad idea. But it is also possible to only distribute a certain amount of free moves (needs to be an even number), say 2 or 4. So for the aforementioned strategy with 10 free moves, it might be best to distribute 4 free moves to attain all sub-strategies containing 6 free moves. In general, in case free moves for white and black are very skewed (say 8+1 free moves), then it is usually better to get rid of the lower number of free moves. In the above case, set `blackPathSplitPieces=all100`, which leads to sub-strategies (8+0) or (x+0) in general.

Modes:

- `whitePathSplitPieces = none` (=default): No such logic is used
- `whitePathSplitPieces = all2`: Distribute max 2 white free move in all possible ways
- `whitePathSplitPieces = all2N`: Distribute max 2N white free move in all possible ways
- `whitePathSplitPieces = all100`: Distribute all white free moves in all possible ways

Same goes for `blackPathSplitPieces`. Btw: The slightly weird parameter name stems from the fact that I have some more ideas of how free move distribution could be defined (e.g. not among all pieces but only a subset of pieces).

These splitting-options are ignored in histogram mode, so that creates the issue, that it is not clear how many induced sub-strategies you get by using these splitting-options. To circumvent this, you can use the new parameter "simulatePlaying" and set it to true (default is false). In that case, the splitting option is applied but the strategies are not played. You can then look at the statistics on the right to see how many sub-strategies you will get (in the problem below, this number is 16)

As an example, the famous problem below by Unto is now solvable in a mere 3 seconds on my machine, when setting "whitePathSplitPieces = all2"

```
Stelvio 4.3
14+16 Stelvio 4.3 Copyright 2023-26 Reto Aschwanden
23.0 Player 1.1/2 00:00:03

r s b q k b s r
P P P P P P P P {P,P,P,P,P,P,P,P} (1/1)
B S B {S,P,P,P,WB,P,P,P} (22/22)
Q P P P P P P
S K

1 Ke1-d1 2 Ke8-e6 0+0
2 Qd1-f3-a3 1 Qd8-d6 14
4 Ra1-a6-d8 1 Ra8-a6
5 Rh1-c8 1 Rh8-h6
2 Bc1-g5-h4 3 Bc8xPe2|e4-c6
3 Bf1-a4 3 Bf8xPa2|a3-e7-f6
Sb1 2 Sb8-d7-b6
3 Sg1-f3-e5-g4 2 Sg8-e7-g6
1 Pa2-a3 (Bf8) 1 Pa7-a5
Pb2 1 Pb7-b5
Pc2 1 Pc7-c5
1 Pd2-d3 1 Pd7-d5
1 Pe2-e4 (Bc8) 1 Pe7-e5
Pf2 1 Pf7-f5
Pg2 1 Pg7-g5
Ph2 1 Ph7-h5
```

Here are some examples where this splitting method turned out to be very useful: [Newly tested problems in PDB](#)

Beware: When using this splitting option, it is possible (albeit very rare), that a shorter solution of same parity might not be found (e.g. a 17.0 move solution when it is a PG in 18.0).

Speculative strategy splitting

With v4.0, a new strategy analysis logic is available in the case there 0 free moves for a given color. In general, Stelvio tries to determine the paths of each piece during the analysis phase as much as

possible. E.g. Ra1x[bSb8]e7 might become Ra1-b1-b7x[bSb8]e7, in case no other option is available. But oftentimes, some options cannot be excluded up front, say Ra1-c4 might be played as Ra1-a4-c4 or Ra1-c1-c4. Before v4.0, this lead to the strategy path Ra1-c4 without further investigation. With speculative splitting available, Stelvio can be advised to split the strategy into 2, one with Ra1-a4-c4 and one with Ra1-c1-c4. With the intermediate square determined in the sub-strategies, it might be the case that one or even both of these options are found to be impossible, since new constraints can possibly be found for the intermediate squares. As you might imagine, all of this is recursive, so in case you have multiple paths like this, say Ra1-c4, Sg1-f4 and Rh1-g5, there is a total of 8 options to add intermediate squares to these paths, so we might get a maximum of 8 sub-strategies because of it.

Modes:

- `speculativeSplitterMode = none`: No such logic is used
- `speculativeSplitterMode = probe`: All 2 move segments (e.g. Ra1-c4) are tentatively split up into sub-strategies at first, but only to determine if the strategy can be further refined or rejected.
 - In case both sub-strategies seem possible after further investigation, then the strategy is not split up at this segment (Ra1-c4 remains).
 - In case only one sub-strategy seems possible after further investigation, then the strategy is refined using the now mandatory intermediate square (e.g. Ra1-c1-c4).
 - In case both sub-strategy are impossible, the strategy is rejected.
- `speculativeSplitterMode = always`: All 2 move segments are split up into sub-strategies (e.g. Ra1-c1-c4 and Ra1-a4-c4).
- `speculativeSplitterMode = alwaysX` (X a positive number, e.g. `always2`): All 2 move segments are split up into sub-strategies, except the first X such segments. So this mode is somewhere in between `always` and `probe`.

Performance impact:

Splitting a strategy into sub-strategies can have both a positive and negative performance impact. My guess is that the mode "probe" is usually better than the mode "always", but I do not have a lot of data to support this. In case strategies are played quickly anyway and a lot of strategies exist (e.g. the massacre SPG case), then "none" is likely the best option here. In any other case, my bet is on the mode "probe".

User interaction

The playing of the current strategy can be stopped by typing 's'. Stelvio will move on to the next strategy thereafter. Pressing Ctrl-C cancels the solving process. In parallel mode, you have different UI pages per seeker/player. You can switch between these pages using left-arrow/right-arrow.

Input / Output

A simple text file (by default `problems.txt`) serves as input. It needs to be in the same directory as `stelvio.bat`. In `problems.txt`, the SPG needs to be given in FEN notation on the first line, and the number of half moves on the second line, something like:

1nbq4/ppk1p3/Rp5p/3npr2/R3P3/2br1B1P/PP2P2P/1NBQNK2

65

Pieces are denoted by:

- K/k = King
- Q/q = Queen
- R/r = rook
- B/b = bishop
- N/n/S/s = knight
- P/p = Pawn

You can add several problems to the input file, which will then be solved in succession.

The result is written into an output file (by default named `problems_out.txt`). There is a `stelvioUI.ini` file for parameters that can be adjusted by the user. If no such file is present, then default values are used. See also `StelvioParameters.pdf`.

Partial solving, strategy conditions

With v4.0, Stelvio can be advised to play only a subset of all found strategies. This can be useful in case full testing is out of reach. It is important to understand that strategies cannot be changed by this functionality. It is e.g. impossible to tell Stelvio that `wSb1` needs to play 6 moves minimum, when Stelvio only sees 4 moves. But what you can do instead is to let Stelvio play only those strategies for which Stelvio has itself calculated at least 6 moves for the `wSb1`. Beware that these restrictions only filter the strategies that are being played. The playing of the strategies that pass the filter remains exactly the same as usual. So even if a strategy condition says e.g. `Sb1<.a` (which means `Sb1` does nothing and remains alive on `b1`), it could still be the case that this knight performs a roundtrip or a switchback (which was simply not detected to be necessary by the strategy finding logic). One other thing to note is that in case of last move retraction due to check in the diagram position, the strategy conditions are applied to the problems that are attained after retracting the possible last moves.

Notation example:

1nbq4/ppk1p3/Rp5p/3npr2/R3P3/2br1B1P/PP2P2P/1NBQNK2

65

StrategyConditions: Pc2-c8=B-e6<.Pf7: AND Ra1-a4<.a

In the example, Stelvio will only play strategies for which: The `wPc2` promotes without capturing on `c8` to bishop and is thereafter captured on `e6` by the `bPf7` AND the `wRa1` moves to `a4` without capturing.

Logical operators:

- AND
- OR

- NOT

These operators are used in the obvious manner. A single clause may not contain both ANDs and ORs. Brackets are needed in this case. E.g. $(A \text{ OR } B \text{ AND } C)$ is not allowed, it needs to be either $((A \text{ OR } B) \text{ AND } C)$ or $(A \text{ OR } (B \text{ AND } C))$. NOT has precedence over AND/OR, as you would expect. So $\text{NOT } A \text{ AND } B$ means $(\text{NOT } A) \text{ AND } B$.

Possible strategy condition types:

- Condition on strategy free moves for a color, e.g.: $fm:w \geq 5$ (i.e. white free moves need to be at least 5). Syntax:
 - Prefix *fm*:
 - Color specifier, either *w* or *b*
 - Numeric condition using one of the operators $<, \leq, =, \geq, >$.
- Condition on capture count for a piece, e.g.: $cc:Sb1 \geq 3$ (i.e. Sb1 needs to capture at least 3 times). Syntax:
 - Prefix *cc*:
 - Piece specifier, something like Sb1 or Rh8
 - Numeric condition using one of the operators $<, \leq, =, \geq, >$.
- Path minimum moves for a piece, e.g.: $ml:Sb1 < 5$ (i.e. Sb1 minimal path is less than 4 moves long). Syntax:
 - Prefix *ml*:
 - Piece specifier, something like Sb1 or Rh8
 - Numeric condition using one of the operators $<, \leq, =, \geq, >$.
- Path condition for a piece, e.g.: $Pa2xSg8 | b8=RxQd8 | c4.a$. Syntax:
 - Prefix: Piece specifier, something like Sb1 or Rh8
 - Followed by 0-n path segments like $xSg8 | b7$
 - See more details in the tables below.

Path segment type example	Meaning
Sb1-c4	Sb1 moves to c4 (does not capture)
Sb1-b1	Sb1 moves to b1 (so switchbacks without capture)
Sb1xb1	Sb1 captures on b1 (so switchbacks with capture)
Sb1xc4	Sb1 captures on c4 (does not capture between b1 and c4)
Sb1_c4	Sb1 moves to or captures on c4 (does not capture between b1 and c4)

Path segment type example	Meaning
Sb1_xc4	Sb1 captures on c4 (and possibly captures 1-n times between b1 and c4)
Sb1__c4	Sb1 moves to or captures on c4 (and possibly captures 1-n times between b1 and c4)

Path endings example	Meaning
Sb1.a	Sb1 is alive
Sb1.d	Sb1 is dead
Sb1<	Sb1 stays on b1 (possibly captured or not)
Sb1<.a	Sb1 stays on b1 and is alive
Sb1<.d	Sb1 stays on b1 and is dead
Sb1<.Qd8:	Sb1 is captured on b1 by Qd8
Sb1-c1<.Qd8:	Sb1 moves to c1 and is captured there by Qd8
Sb1.Qd8:	Sb1 is captured somewhere by Qd8

Capture example	Meaning
Pa2xb5	Pa2 captures some piece on b5 (as next capture)
Pa2xQd8	Pa2 captures Qd8 somewhere (as next capture)
Pa2xQd8 b5	Pa2 captures Qd8 on b5 (as next capture)

Promotion example	Meaning
Pa2=R	Pa2 promotes somewhere to rook
Pa2-a8=R	Pa2 promotes on a8 to rook
Pa2xb5=R	Pa2 captures on b5 and goes on to promote somewhere to rook
Pa2xb5=R-c4<	Pa2 captures on b5 and goes on to promote somewhere to rook and ends on c4.

Castling example	Meaning
Kg1-g1-h2<	White king side castling, then end on h2
Kc1-c1-h2<	White queen side castling, then end on h2
Kg8-g8-h2<	Black king side castling, then end on h2
Kc8-c8-h2<	Black queen side castling, then end on h2

Read/write strategies to file

Stelvio can be advised to save all found strategies in structured format to disk. In a second step,

instead of searching for strategies, Stelvio can thereafter read these strategies from disk and try to play them. This is especially useful in case one uses histogram mode at first. Calculating the histogram sometimes requires a lot of time, only to find very few strategies. In order to subsequently play these strategies, they can now be read from disk in no time, instead of recalculating all of them all over again. Limitation: Last move retraction and save/read strategies from/to file are not currently supported in combination.

Known issues

Given a cooked SPG and a cook-strategy s for it (i.e. a strategy s that can be played in multiple ways and therefore cooks the SPG). It is possible that not all permutations of how this strategy s can be played are found. This should usually be irrelevant, as the verdict is unaffected (the SPG is cooked, the strategy s is found to be a cook-strategy in any case).

Explicitly finding all the different ways a cook-strategy can be played seems pretty pointless, so I have not invested time in fixing this. As possible fixes imply a performance penalty, I'm inclined to leave this as is for the moment at least.